Communicating Risk in Public Pension Plans

Robert T. McCrory, FSA

Introduction

Over the past few years, actuaries have been discussing how best to communicate risk to their clients. Much of this debate has been focused on the public sector, in which employers and retirement boards composed largely of laypeople make decisions involving billions of dollars of benefits and taxes.

This article aims to present a number of approaches to communicating and discussing risk with employer and retirement board members. We mean to be practical here: All of these techniques have been used in client situations, and each of them is useful within the limitations discussed.

This paper is a written version of a presentation at the Employee Benefits Meeting sponsored by the Society of Actuaries on June 5, 2008. Anyone who is interested can obtain the PowerPoint presentation from the Society's Web site, [*www.soa.*org](http://www.soa.org), or from me directly.

Our Approach

We will review and score six different approaches to communicating risk. We should take a moment to discuss the nature of the risk we are trying to communicate. In general, we are trying to roughly quantify and convey the routine market risk faced by a public pension plan. We will measure this as the expected variability in employer costs and funding ratio during "normal" financial markets.

It is an altogether different issue to quantify and discuss the *maximum* risk or variability faced by a public plan. This requires a much different approach—to the extent that it is even possible—and we will not be addressing that issue here.

We will review the following vehicles for discussing risk with clients. The order of presentation is in my order of preference, as will become clear as we proceed.

- Full Simulation
- Partial Simulation
- Error Bars
- Back-testing
- Range Projection
- Simple Projection
- Risk-Free Discount Rate

In discussing these approaches to risk, we will evaluate each on a number of criteria, as follows:

- Simplicity

  We want our measure of risk to be as simple as possible, while still communicating the nature and extent of risk fully and accurately. This will involve some tradeoffs, as we shall see.

- Relevance

  The risk measured should be relevant to our clients. This will usually be a matter of degree and of subjective judgment.

- Probative Value

  A risk measure can reveal interesting facets of pension plan operation and dynamics. This represents a great opportunity for learning about our plans and passing this knowledge along to our clients.

- Consistency

  We would like our risk measures to be consistent from year to year. Changes in the measure should reflect changes in the plan, not merely changes in the economic environment.

- Cost

  Calculation and presentation of risk measures is as important as the calculation of plan liabilities and costs. Ideally, we would want these risk measures to be produced as a routine part of all actuarial studies, with as little impact on our internal costs and fees as possible.

- Clarity

  Any measure of risk should be understandable to lay people: board members, plan staff, elected officials, voters and the press. Above all, our risk measures should not mislead.

Each approach will be scored as “good,” “fair” or “poor” for each of the above criteria. The scoring is my own, and is necessarily subjective. Your assessment will almost certainly differ in some areas—maybe in all areas—from mine. Good. The methods covered here will have differing usefulness based on each actuary’s experience, capabilities and client base.

Let’s proceed.

Full Simulation

In a full simulation of a pension plan, we build a complete computer model of all aspects of the plan worthy of study. Such a model will include all of the member data, benefit provisions, and actuarial assumptions that would be present in a typical pension actuarial model. However, for our purposes here we must go further, and include elements not usually present in computer systems designed to produce actuarial liabilities and costs. Additional elements will include at least some of the following:

- Projection of active and inactive member populations to future years. The projection logic will provide for new plan members necessary to achieve stability, growth or decline in the active census. Furthermore, in systems with multiple benefit tiers, inter-population dynamics will be programmed, resulting in terminations or retirements in an older tier triggering new hires in the new tier or tiers.

- Stochastic simulation of future asset returns and inflation.

- Complete simulation of special benefits that depend on variability. For example, benefits that depend on “excess” investment earnings over the actuarial assumption should be modeled in full detail.

- Full modeling of the actuarial valuation process, to allow the computation of future actuarial liabilities and costs based on simulated results. The model will include the creation and amortization of accrued liability bases, and the calculation of the actuarial value of assets. In addition, if the process of computing actuarial costs is subject to legal constraints, such as a limit on the amount by which actuarial contributions can increase, such constraints should be included.

- Depending on the circumstances, other issues may be included in the model. For example, if the funding of the plan depends on property tax revenues, it may be appropriate to model the relevant property tax base to assess the required contributions relative to available revenue.

The purpose of such an elaborate model is to understand the behavior of the plan. Note that we are interested in the *dynamics* of the plan—how it behaves over time, how it responds to changes in its environment—rather than the status of the plan at a single point in time. Of particular concern is understanding, quantifying, and communicating the risks associated with the plan, especially those having to do with cost.

Consider an example. Graph 1 below shows the simulated actuarial cost for 30 years for a group of statewide plans. In this system, there are five plans: one each for state employees, local employees, school employees (other than teachers), state safety workers, and judges. Assets totaled $12 billion. For purposes of our study, the plans were grouped into a higher funded group (state employees, safety, and judges, shown in the graph) and a lower-funded group (local and schools).

There are a few points worthy of mention in Graph 1. First, look at the variation (individual trials are shown in blue). The current cost of the combined plans is about 10 percent of payroll, but there are trials in which the cost in nearly quadruple that figure within 30 years.

Second, even if we ignore the extremes, we note that the quartiles (shown in green), range from 0 percent to 20 percent of payroll. Therefore, based on this asset model, we can say with 50 percent confidence that the cost will range from nothing to twice current levels. This pattern—quartiles of 0 percent or double current cost—has occurred repeatedly with many different clients in our simulation work. It has become a tentative rule of thumb to use for advising clients, even in the absence of such a study. We will have more to say about asset models later.

Third, Graph 1 is a powerful tool for communicating the cost risk of a pension plan. Sometimes it can be a bit of a shock to a board, but most of the time it confirms what they already knew: We're guessing at the cost almost as much as they are.

Graph 1

Simulated actuarial cost for a group of statewide retirement plans.
Blue lines show individual trials; 50 of 500 trials are shown.
The red line is the annual average of all trials, and quartiles are shown in green.

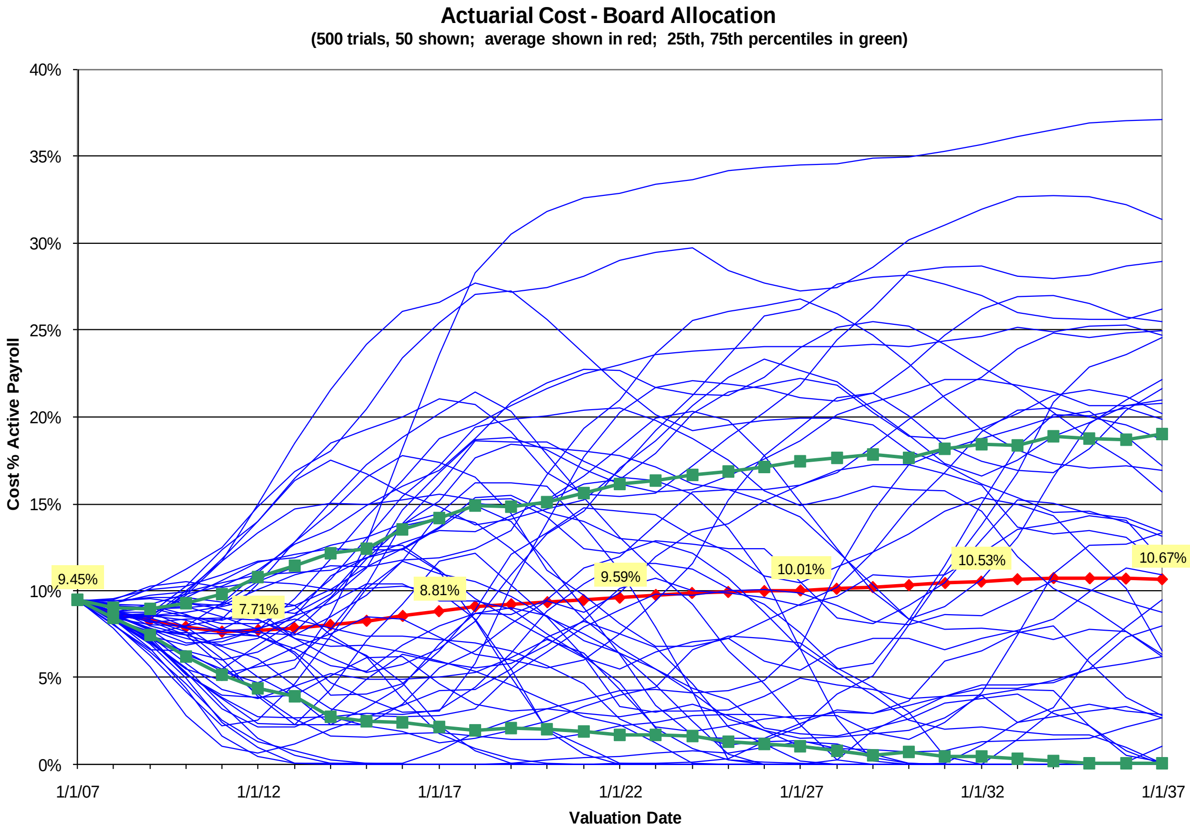

A simulation of actuarial cost necessarily involves simulating the funding ratio, the ratio of the actuarial value of assets divided by the actuarial accrued liability. The simulated funding ratio for the higher funded group of this system is shown below in Graph 2. The format of this graph is the same as Graph 1: 50 of 500 trials shown in blue, annual average funding ratio in red, quartiles in green.

Graph 2 reveals a number of interesting features. First, notice that the trend of the mean funding ratio is upward, exceeding 100 percent in about 15 years. This occurs because the funding ratio graph is skewed upward: Quite a few trials exhibit runaway overfunding, pulling the average funding ratio up.

Second, note that the funding ratio seldom drops below about 50 percent. In fact, in most plans the funding ratio *never* drops below 50 percent. This floor on the funding ratio occurs because when assets are that low, the actuarial cost is quite high, higher than the sum of benefits and the maximum potential market loss on the reduced plan assets. This guarantees positive cash flow and an improvement in funding. Therefore, while overfunding can occur, underfunding is strictly limited by the dynamics of actuarial funding.

Third, note that in this example, the funding ratio indeed drops below 50 percent; in fact, if we simulate further in time, the funding ratio in some trials reaches 0 percent—loss of all assets—after a number of years. The reason for this is that there is a statutory limit on the amount by which employer contributions can increase; the law limits increases to no more than 0.6 percent of payroll per year. Consequently, in some trials the actuarial losses accumulate, and the contributions can't increase fast enough to compensate, causing the plan to become insolvent.

This kind of dynamic behavior is impossible to see without a simulation. In fact, there is more to this cost limitation than meets the eye. This leads us to a digression.

Graph 2

Simulated actuarial funding ratio for a group of statewide retirement plans.
Blue lines show individual trials; 50 of 500 trials are shown.
The red line is the annual average of all trials, and quartiles are shown in green.

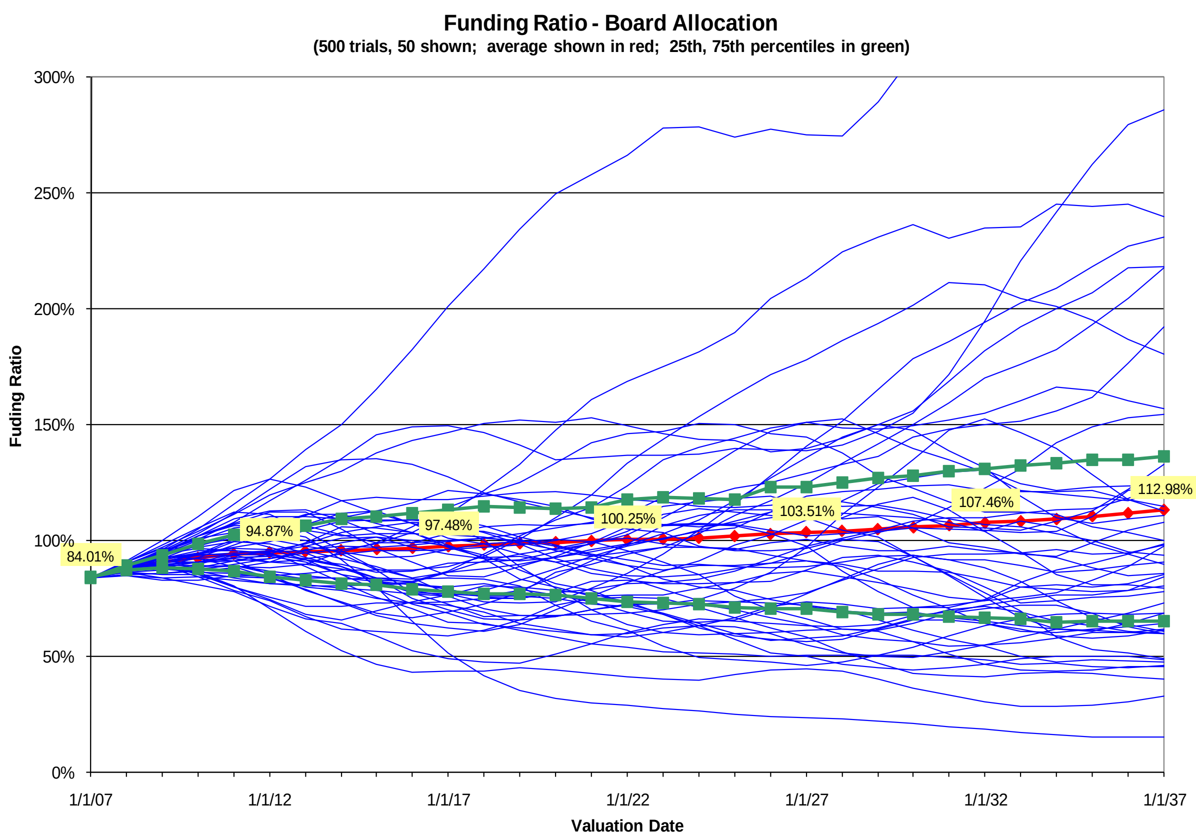

Ruminations on a Smile

Consider Graph 3 below. It looks like a smile, doesn't it?

Graph 3 shows the funding ratio produced by one simulation trial in our statewide plans. The higher-funded plan group, the subject of Graphs 1 and 2, is shown in red, while the lower funded plan group is shown in blue. The same series of random economic results—investment returns and inflation—was applied to each plan group, but we note that the resulting behavior is very different.

The funding ratio of the higher funded plans starts at over 80 percent, and with favorable returns increases to over 100 percent within 10 years. Then the markets turn around and the higher funded plans become insolvent. On the other hand, the lower funded group starts out at a bit over 60 percent funded, improves, then drops with poor returns, but rather than becoming insolvent the lower funded plans eventually become over funded, with funding ratios reaching 150 percent, even as the higher funded plans are losing all of their assets.

It is important to note that these plans are not very different. Other than their starting funding ratios, they are nearly identical. Both groups of plans contain large groups of state workers. Benefit provisions are nearly identical. Indeed, in most of the simulation trials these plans have very similar results.

Graph 3

One simulated trial of the actuarial funding ratio for two groups of statewide retirement plans. The blue line is the funding ratio for the lower funded group, the red line for the higher funded group. The same economic results – returns and inflation – were applied to each plan group.

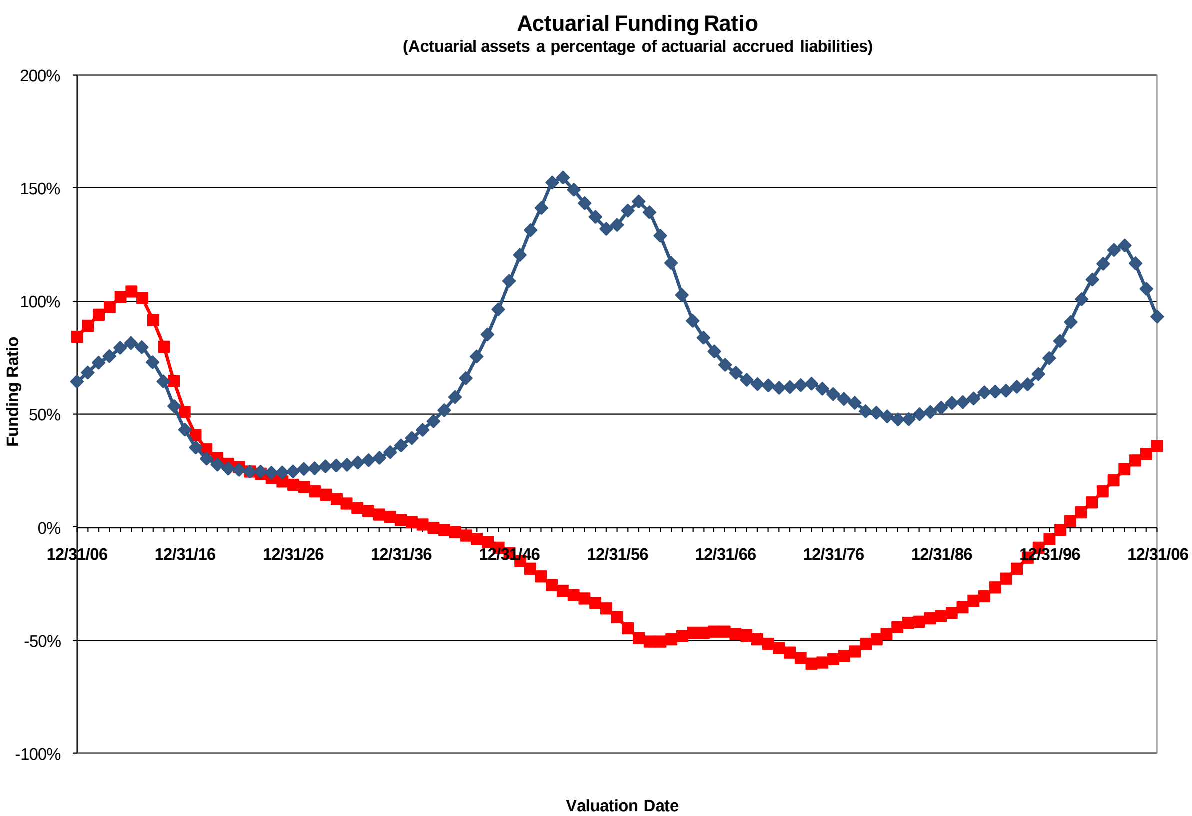

So, what's going on here? Remember earlier that we mentioned that there was a cap on increases in the employer contribution? The employer contribution can only increase by up to 0.6

percent of pay per year, regardless of the change in the actuarial cost. It is this cost cap that causes the puzzling behavior in Graph 3.

Oddly enough, the higher funded plans find themselves at a disadvantage. Both sets of plans gain assets with the favorable returns in the first five years or so. Then, when the market becomes unfavorable, the higher funded plans are at a double disadvantage: Their assets are higher, so they lose more money in the markets, and their actuarial cost is lower, so the cost cap keeps employer contributions lower for a longer time. The result is that the higher funded plans never catch up to the market.

This is mathematically chaotic behavior: extreme sensitivity to initial conditions. It is caused by the limit on annual employer contribution increases. Who knew?

This reveals a huge benefit of simulation: the ability to test seemingly innocuous policy options and discover their full implications.

Full Simulation: Evaluation

This leads us to a discussion of the advantages and disadvantages of a full simulation based on our list of evaluation criteria.

- Simplicity – Fair

  As my colleague Greg Stump says, sometimes the answer is not a number. The results of a simulation study are not simple: There is no single result, or "real" value. How could there be? The future is not simple. Risk isn't simple.

  A simulation study is a complicated undertaking that produces complicated results. Communication of these results requires considerable planning and effort. On the other hand, there aren't a lot of caveats and disclaimers that need to be made, such as "if all actuarial assumptions are realized." The construction of a comprehensive model makes most of these unnecessary.

- Relevance – Good

  A simulation study is flexible and is therefore capable of producing many measures of performance and risk. In our studies we have modeled such items as funding under statutory cost limitations and plan cost in comparison with a property tax levy. Accordingly, a simulation study can be made as relevant as we have the knowledge, time and energy to make it.

- Probative Value – Good

  As noted above, seemingly innocuous plan provisions or funding methods can have unexpected results. The surprising result of a cost limitation—chaos—was presented above. But there is a long list of suspects in this regard: Investment-driven benefits, benefit floors, and sharing "excess" earnings with plan members are only a few.

As another example, we have found that boards usually have an exaggerated notion of the impact of actuarial asset smoothing on the volatility of their plan contributions. Nearly every time we present a cost simulation, the first question asked is "What about actuarial smoothing?" When it's already in the simulation model, the answer is: "It's already there."

- Consistency – Good

  We have found that simulation studies show excellent consistency from year to year. The economic model used in the simulation model may be adjusted from year to year to reflect changes in the long-term economic outlook, but short-term market fluctuations will not usually be reflected. Accordingly, while the current starting values for simulated plan costs and liabilities will change, the overall distribution and dispersion of the simulated results will not vary significantly.

- Cost – Poor

  A full simulation requires a computer system that can be economically programmed, and the time and resources to do so. This can be daunting for both small and large actuarial firms. Small firms often lack the resources, and large firms with the resources must often charge prohibitive fees for their work. In addition, most actuarial systems are designed for the computation of present values, not for the projection of cash flows, liabilities and actuarial costs.

  Therefore, full simulation can be costly, perhaps prohibitively so, unless the actuarial modeling system is designed from the outset to facilitate a seamless introduction of simulation into routine annual actuarial work.

- Clarity – Good

  We have found that simulation results are easy to present to lay boards, especially when the graphics are clear and readable. The display of possible futures, in particular the potential variation in plan cost, is clear and powerful. These graphs are vastly easier to understand than text or statistical measures.

  However, a word of warning: Most boards are fairly sophisticated graphically, and we have been scolded from time to time for graphs that were small, hard to read, or labeled unclearly. Take some time with your graphics.

  Insofar as communication with the press and taxpayers is concerned, our experience has been quite favorable. The chairperson of one of our boards is a leader of a major anti-taxation group; he has been quite supportive as we have used simulation approaches to explore contribution risk. In addition, simulation studies do not present a single number or result that is easily misinterpreted in the press.

It is with particular reference to cost that we consider partial simulation.

Partial Simulation

Full simulation requires capabilities not usually found in traditional actuarial software packages. Fortunately, there is an alternative approach that achieves most of the goals of a full simulation but uses traditional software.

Most actuarial software packages are capable of projecting plan liabilities and normal costs into the future. Such projections are usually made assuming the plan's actuarial assumptions are exactly met, so no variation in economic scenarios is allowed for. However, economic uncertainty usually affects future liabilities only through the impact of inflation on payroll and benefits, which in turn has a much smaller impact on future plan costs and funding ratios than the performance of plan assets. This suggests the following approach:

1. Use existing actuarial software to project future benefits, payroll, normal costs and liabilities as if all actuarial assumptions are met exactly.

2. Transfer the above projections to a spreadsheet. Program the spreadsheet to compute future assets and plan costs using the plan's funding method and the projected payroll, benefits and liabilities. Add a column for each future year's return on plan assets; for initial programming, make every year's return the actuarial assumption, but keep the spreadsheet flexible enough to vary the asset returns in any year. Add programming for the actuarial value of assets.

3. Derive and check deterministic cost projections.

4. Replace the annual return with a simulated return for each year, using the asset model of your choice. Collect, analyze and graph the resulting simulated costs and funding ratios.

Graph 4

Simulated actuarial cost for a large county retirement plan.
Gray and black lines show individual trials; 50 of 500 trials are shown.
The orange line is the annual average of all trials, and quartiles are shown in blue and green.

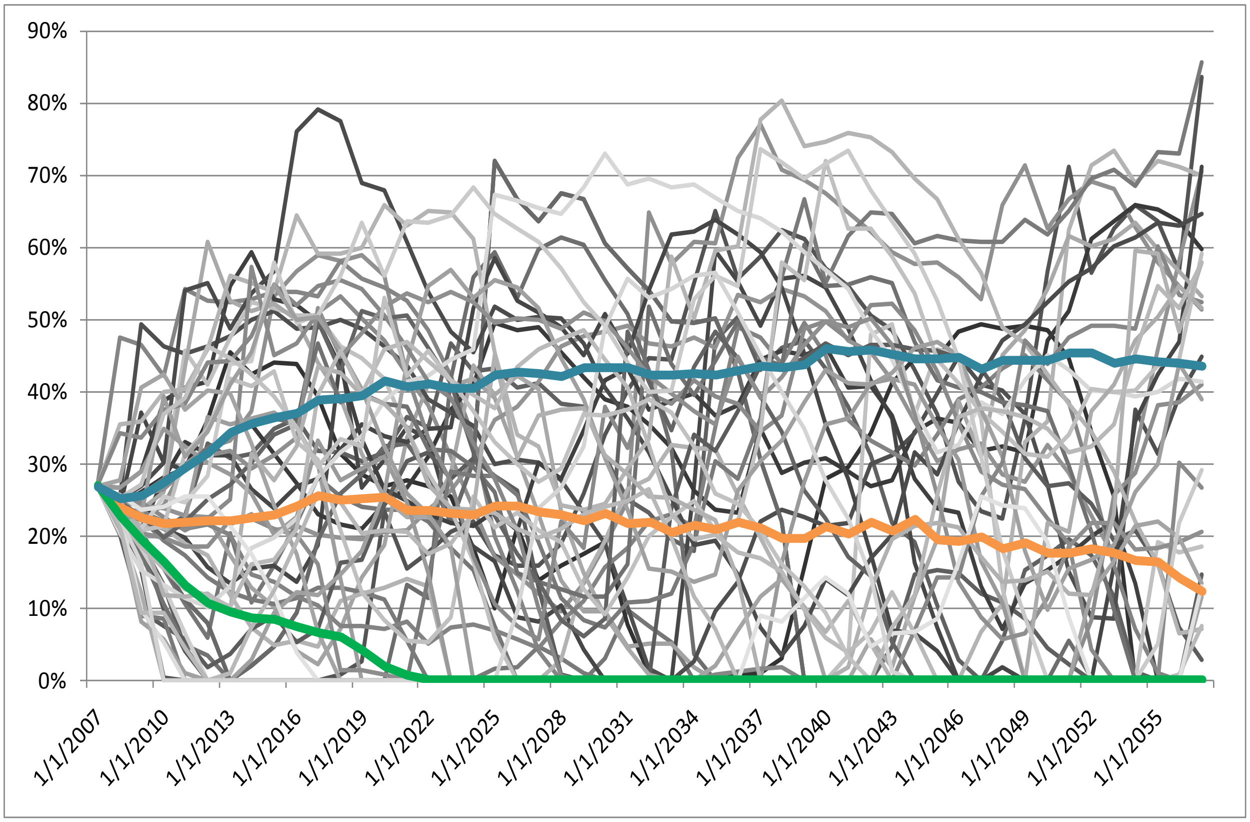

Graph 4 above shows the result of such a partial simulation study.

This approach is designed for practicality: Simulation results are available as an extension of the basic actuarial valuation with a minimum of additional effort. To be sure, we have sacrificed some accuracy. In particular:

- Demographic assumptions are assumed to be met exactly. This is not much of a sacrifice. Demographic gains and losses add very little to cost volatility in larger plans unless they result from structural changes in the workforce – layoffs, mergers, work rule changes in union negotiations and the like, which are not usually included in our models anyway.

- The largest source of variation omitted from this approach is that due to inflation. Clearly, inflation causes changes in payroll, benefits, liabilities and cost that, depending on the plan, can be significant.

Nevertheless, a partial simulation can be a valuable extension to an annual actuarial valuation. At a minimum it gives the plan board, staff and stakeholders some indication of the extent to which future plan costs may vary from current levels. That alone is enough to justify the effort involved.

In terms of our criteria for rating methods of communicating risk, the same points made for full simulation continue to hold, albeit with a few caveats.

- Simplicity – Fair

  Presenting the results of a partial simulation requires the same degree of care and planning as a full simulation study. However, the simplification involved in fixing inflation at a single level should be disclosed and discussed.

- Relevance – Fair

  A partial simulation study can usually investigate the same items as a full simulation study, so its degree of relevance is about the same. However, there is a notable exception: Plans with investment-driven cost of living increases have benefits, and hence liabilities, that can be highly dependent on inflation. For this category of plans, and for some others, a partial simulation study is probably inadequate, and may even be misleading.

- Probative Value – Good

  In most cases, partial simulation studies have considerable probative value. Again, however, note that partial simulation is not appropriate for all plans, and for some plans may be incapable of quantifying key sources of risk.

- Consistency – Good

  Generally, partial simulation analyses have the same consistency from study to study as full simulations.

- Cost – Good

  In most cases, a partial simulation represents a minor addition to the effort required in an annual valuation. In my view, the additional information provided to the client is well worth the minor expenditure of effort.

- Clarity – Good

  The results of a full or partial simulation study can be presented in about the same way and are equally understandable to boards and staff.

Therefore, a partial simulation study offers us the opportunity to include a comprehensive discussion of risk in our presentation of annual actuarial results, with a minor extra effort on our part.

Before discussing other approaches, I feel another digression coming on.

Fear of Flying: Asset Models

Don't get too hung up on asset models.

In preparing a simulation or partial simulation, you will have to develop a stochastic model to produce investment returns and inflation rates. A lot of ink has been spilled on the pros and cons of various solutions to this problem. Much of it has been wasted.

In the community of computer modeling and simulation, one phrase is paramount: "precision for decision." The model you build should only be precise enough to help you make the decisions you plan to use it for. Additional precision and effort is not only unnecessary, but often downright harmful.

Our goal in the communications discussed here is to give decision makers an idea of the volatility in plan liabilities and costs that is to be expected routinely. We are not interested in extreme events, at least not in this context. Consequently, use a simple return model; a normal or lognormal distribution of returns and inflation with reasonable means and standard deviations is fine for our purposes. Yes, I know that the actual tails are fatter and the distributions aren't symmetric (or even stable), but we really don't care for the *basic* communication of risk.

If we were developing an asset allocation study or trying to determine the worst possible case, more care would be required in selecting an asset model. Above all, don't let the perfect drive out the good. Don't allow a search for excellence in an asset model stop you from communicating the basic dimensions of risk to your clients.

Error Bars

I recently read a popular science book that described attempts to determine the age of the universe. The book was written before the age was known with the precision determined today (13.7 billion years), so there was considerable uncertainty in the various estimates. What I discovered in reading the book was that the researchers spent as much time determining the range of error in their estimates as they did making the estimates in the first place. We can learn something from that.

We spend a great deal of effort computing liabilities and costs; can't we spend a little effort determining the error bars around our computations? One possible approach would be to use simulation results to find the variability in liability or cost computations. The table below shows how cost information could be presented.

TABLE 1
Range of actuarial costs for a large county retirement plan
Median and quartile costs are shown both now and in the future.

	0 Years (Valuation Date)	5 Years	10 Years	20 Years
75th Percentile	**26.86%**	**31.53%**	**38.85%**	**43.54%**
Median	**26.86%**	**21.95%**	**25.49%**	**21.00%**
25th Percentile	26.86%	10.49%	6.54%	0.00%

In Table 1 we see the computed actuarial cost on the valuation date – Time 0. As time goes on, the median contribution declines as the actuarial accrued liability is amortized. However, uncertainties in investment returns and inflation increase the degree of uncertainty as time progresses.

The information in the above table could be displayed a little differently. In Table 2 below, the quartiles are shown as additions or subtractions from the current cost; the median cost is not shown.

TABLE 2
Range of actuarial costs for a large county retirement plan
Future quartile costs are shown relative to the current cost.

	0 Years (Valuation Date)	5 Years	10 Years	20 Years
75th Percentile		+4.67%	+11.99%	+16.68%
Current Cost	26.86%			
25th Percentile		-16.37%	-20.32%	-26.86%

A simulation is not necessary to show the range of costs. A simpler approach could be used, as follows.

1. Determine a reasonable distribution of real returns. As noted above, as long as the mean and variance are reasonable, in most cases the details of the actual distribution don't matter much.
2. Using traditional actuarial methods, compute liabilities, cost and funding ratio for a reasonable sample of these real returns, using the assumed inflation rate.
3. Weight the costs and funding ratios by the probabilities determined in the first step, and determine a distribution of costs.

This approach was used to compute a distribution of actuarial costs as shown in Graph 5.

Graph 5

Distribution of actuarial cost for a transit district plan.

Based on the current set of actuarial assumptions, the actuarial cost was computed at 10 percent of pay. The actual cost will depend on future investment returns; the distribution of the unknown actual cost is shown below. The horizontal axis is the actuarial cost as a percentage of pay. The vertical axis is the number of trials in which the actual cost was realized.

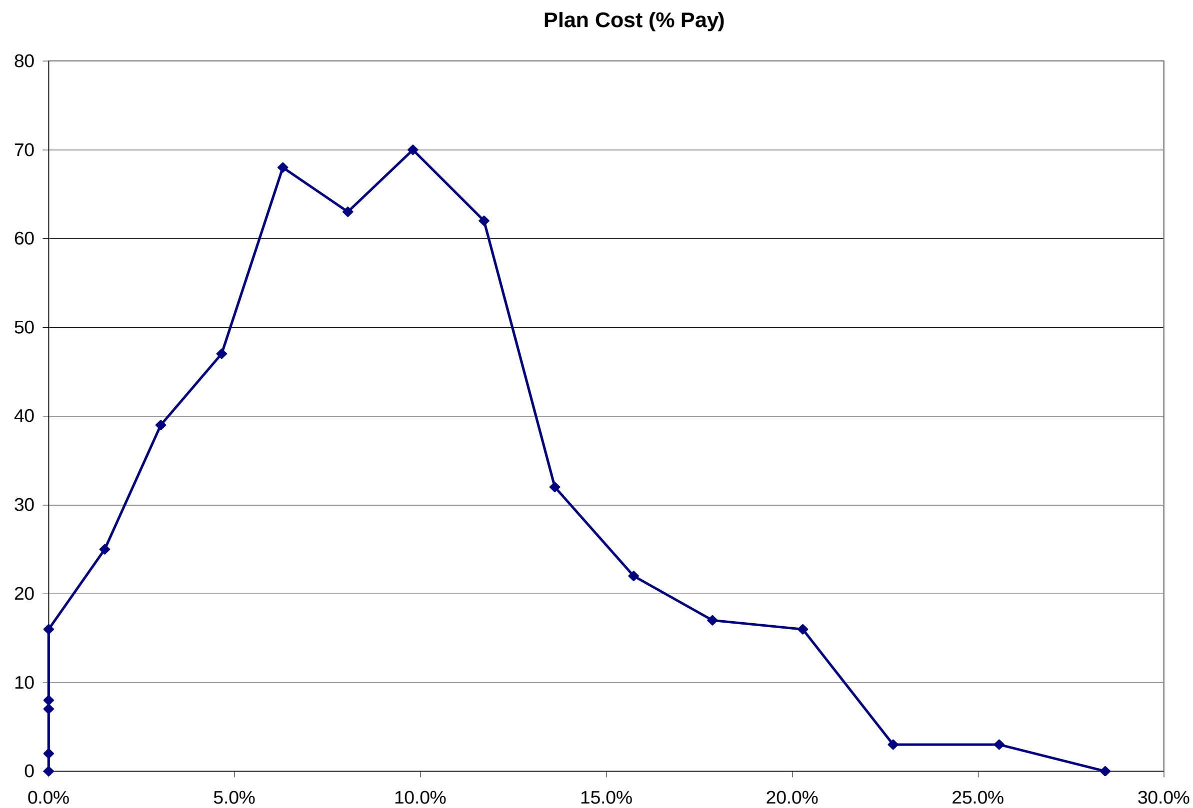

In terms of our criteria for rating methods of communicating risk, we can say the following about error bars.

- Simplicity – Fair

  Showing our costs and funding ratios, plus or minus an error factor, is a simple and concise way to express the uncertainty in our estimates. Whether presented graphically or numerically, it should lead to a focused and useful discussion with our clients.

  It may be that some clients will be surprised that there is any error bar, any degree of uncertainty at all. They are the ones who most need the discussion.

- Relevance – Fair

  A full understanding of the uncertainty of current cost and funding estimates is certainly relevant to our clients. However, there may be other important items, such as the ability of an existing tax base to support plan contributions, which may require a full or partial simulation to quantify.

- Probative Value – Poor

  Error bars are certainly useful, but they do not allow the exploration of the future dynamics of a plan. Are costs tending higher or lower? Is the uncertainty in plan costs increasing with time? In short, is the plan risk increasing? A full or partial simulation is usually necessary to answer these types of questions.

- Consistency – Good

  Error bars remain fairly consistent relative to cost or funding ratio from year to year, so consistency is usually not a problem.

- Cost – Good

  Computing the range of error in our estimates is not a significant addition to the effort required in an annual valuation. Moreover, it seems to be professionally responsible to disclose this type of information.

- Clarity – Poor

  Depending on the board and staff, error bars may be more or less understandable than a simulation. A great deal depends on the learning style of the board and their familiarity with quantifying and disclosing errors. For example, public water districts usually employ graduate engineers, and they find their way onto pension boards and into executive positions. In these cases, they will be very familiar with error bars and little explanation will be necessary.

  On the other hand, members of the public and the press are likely to find error bars a difficult concept to fathom. Even the term itself may cause alarm.

Backtesting

Backtesting is a “what if?” analysis in retrospect: What would plan costs and funding ratios look like over the next 10 years if the economy exactly repeated the last 10 years? Obviously, the 10-year period cited here is arbitrary—it could be two, five or 20, depending on the point you want to make. It does not need to be the last *n* years: You may choose any period you want, but you should be prepared to defend your choice.

Backtesting is simple: Simple to do and simple to explain. A modest amount of work on a spreadsheet will accomplish our goals. Moreover, it is frequently informative, and it can serve to remind our clients that variation in cost and funding ratios is an expected and unavoidable part of planning.

Graph 6 below shows the projection of plan costs for a transit district assuming the repetition of the investment experience of the prior four years. In fact, during 2000, 2001, 2002, and 2003, the return on Plan assets was 0.39 percent, 2.91 percent, -25.78 percent, and 24.44 percent, respectively. Note how the full realization of the investment returns is delayed by the actuarial smoothing of assets. After the four-year backtesting period, the investment return was assumed to be the actuarial assumption.

Graph 6
Projection of actuarial cost for a transit district plan.
The cost of the plan was projected as if the last four years of investment experience were to repeat themselves for the next four years. The horizontal axis is years. The vertical axis is cost as a percentage of pay.

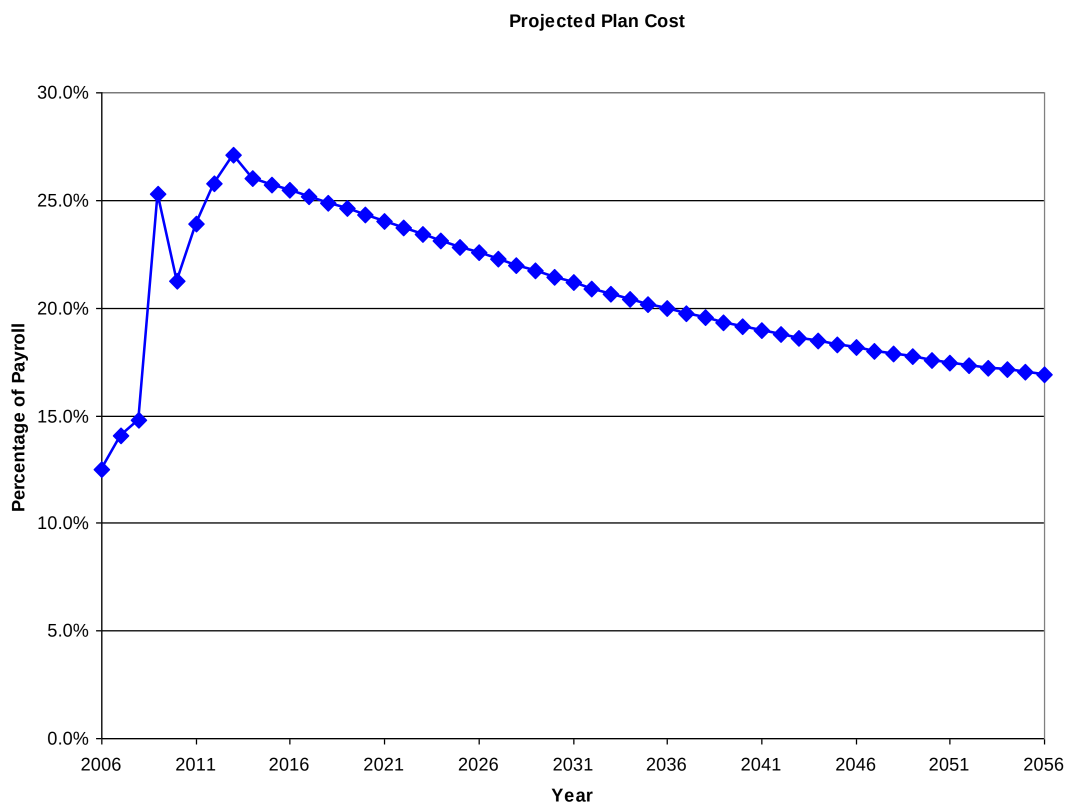

This projection served as a reminder to the board that the cost of the plan could more than double within a few years. In particular, a benefit improvement was being considered, and the reminder was timely. In terms of our criteria for evaluation:

- Simplicity – Good

  As noted above, backtesting is possibly the simplest technique we will discuss.

- Relevance – Fair

  The future performance of the plan is certainly relevant to policymakers and stakeholders. Reminding these groups of the possible variation in cost and funding ratio is always a good idea. However, the last few years will not repeat themselves, so the value of this projection is fairly limited.

- Probative Value – Poor

  A single projection has some value but, as in the case of error bars, the dynamic behavior of the plan is neither explored nor revealed. A full or partial simulation is usually necessary to investigate these issues.

- Consistency – Poor

  The results of backtesting will change over time if different time periods are chosen for the past returns. Therefore, the backtesting projection is likely to look significantly different each year.

- Cost – Good

  Backtesting is an easy and quick projection on a spreadsheet.

- Clarity – Good

  It is difficult to misinterpret a backtesting projection. It is clear what it is, and what it is not. I am not sure how the public and the press would interpret the results of backtesting; in particular, I do not know if they would find it instructive or relevant.

Range Projection

It seems that the most commonly used risk communication technique is the range projection. The idea is simple: Project future costs assuming all assumptions are met, then with say, 1 percent added or subtracted from the return on assets. The result is usually something like Graph 7.

Note that this is the same plan as was shown in the partial simulation in Graph 4. Referring to Graph 4, note that the range of costs ranges from 0 percent to 80 percent of payroll, more than twice as broad as the range of costs shown in Graph 7.

It is in this area that range projection fails. The full scope of the risk faced by the plan is hidden. Indeed, there is a strong visual tendency to focus on the area between the outer projections, assuming that the actual cost will be somewhere in between. In actuality, a simulation study will show that many results are well outside the narrow range shown using this technique. Broadening the range doesn't help much: The viewer still focuses on the area between the projections, and the idea that possibilities lie well outside of this area tends to be lost.

Accordingly, a range projection must be used cautiously, and explained thoroughly, to have the communication value we want. This is reflected in our scoring.

- Simplicity – Good

  A range projection is a simple variation on a traditional actuarial projection. It is widely in use. It is easy to explain the basics but, as noted above, there is hidden risk that must be pointed out.

- Relevance – Fair

  Some indication of the range of future outcomes is clearly relevant to policymakers and stakeholders. However, the lack of probabilistic variation limits the usefulness of this approach.

Graph 7
Projection of actuarial cost for a large county plan.
The cost of the plan was projected at the current actuarial rate, then at that rate plus or minus 1 percent per year. This is the same plan as was shown as a partial simulation in Graph 4. The horizontal axis is years. The vertical axis is cost as a percentage of pay.

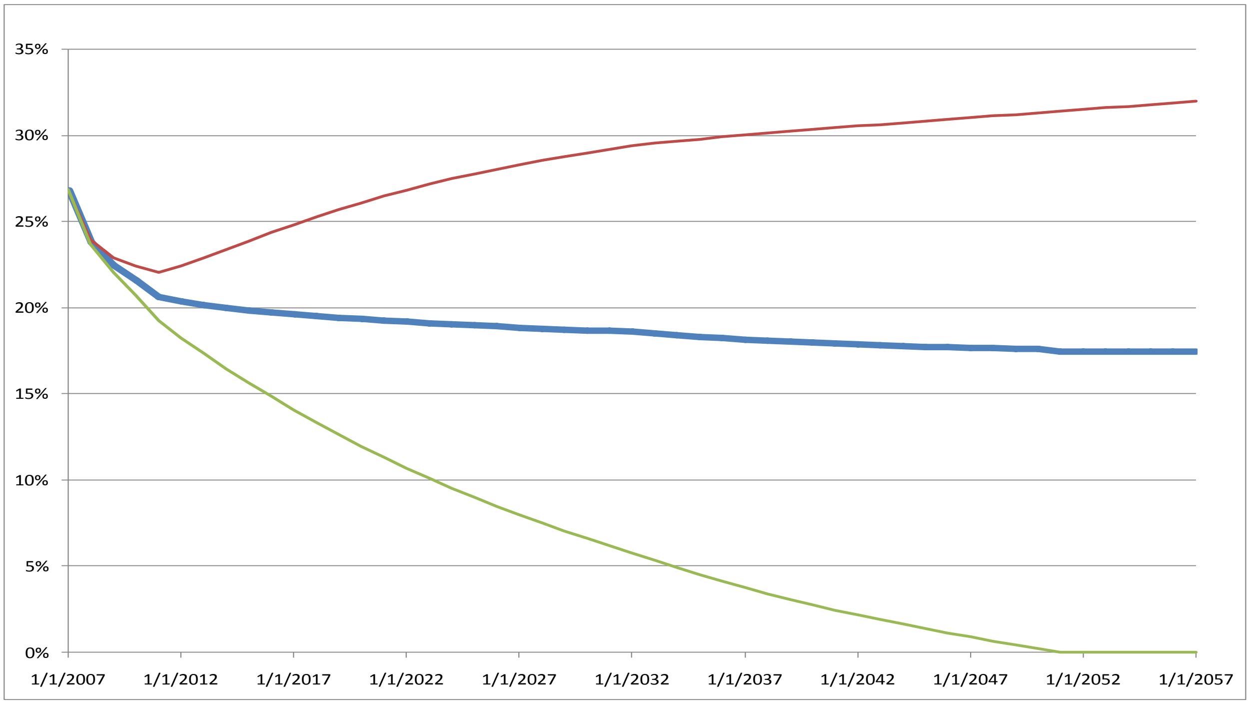

- Probative Value – Poor

  A range projection has some value but, as in the case of backtesting and error bars, the dynamic behavior of the plan is not part of the study.

- Consistency – Good

  Other than the starting point and the delaying effects of asset smoothing, there will be little annual change in a range projection.

- Cost – Good

  This is a quick and easy technique to use: Another spreadsheet application.

- Clarity – Poor

  Generally, the range of outcomes shown in range projection is much too narrow. Moreover, the eye tends to fill in the space between the projected results with all future results, ignoring the significant likelihood that results will be literally off the chart. Accordingly, even with verbal caveats, this technique can be seriously misleading.

Simple Projections

The next technique sometimes doesn't communicate risk at all, but it still should be a routine part of our annual valuation work. Why? We often tell our clients that the actuarial cost will remain level if the actuarial assumptions are met. But this is almost never true. Take a look at Graph 8 below.

In Graph 8 we see the progression of the actuarial cost of a fully funded plan if all actuarial assumptions are met exactly. The projection is for 50 years, during which time new members are assumed to join the plan to replace those members who retire or otherwise terminate. The new members assumed are based on the age profile of recent hires. We note that the cost trends upward: Why?

After some research, we determined that over the past 20 years or so, the average age of a newly hired transit operator had increased from about age 32 to nearly age 40. As these older hires replace members hired at younger ages, the normal cost of the plan increases. We find many other public plans in this situation.

Had we not done this simple projection, we would have proceeded from year to year, generating small losses and cost increases, without really knowing why—and without being able to communicate the expectation of future cost increases to our client.

A simple cost projection can show underlying trends. We've seen projected cost changes from the operation of asset smoothing, from the operation of the funding method, and from new tiers of benefits phasing in.

Graph 8
Projection of actuarial cost for a transit district plan.
The cost of the plan was projected assuming all actuarial assumptions are exactly met.
The plan was approximately fully funded.
The horizontal axis is years. The vertical axis is cost as a percentage of pay.

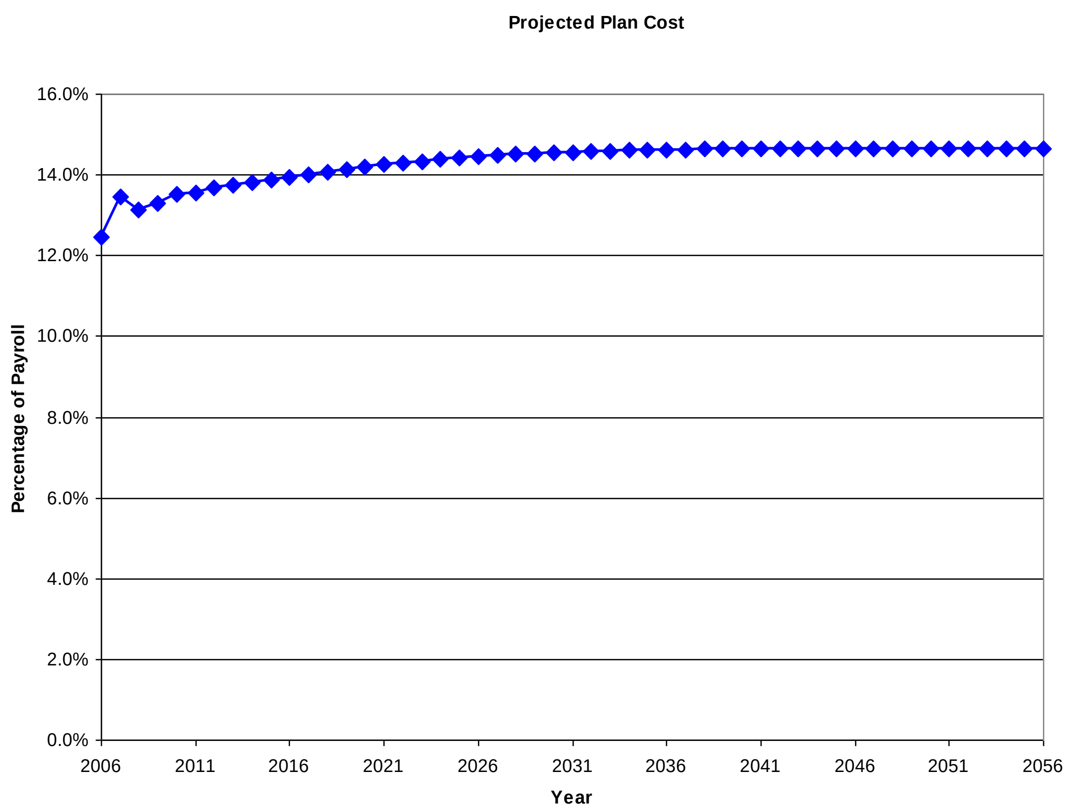

However, there is another type of simple projection that can tell us a great deal about risk: This is the projection of future assets as a percentage of payroll, an example of which is in Graph 9.

In Graph 9 we see that the ratio of plan assets to covered payroll increases from about 5.5 times on the valuation date to about 9.5 times in 50 years. This increase occurs as the funding of the plan improves (it is currently about 90-percent funded), but that is a minor influence. Assets grow as the workforce matures and as improvements in benefits enacted a few years ago produce higher liabilities and benefits. Taken together, these influences cause the ratio of assets to payroll to nearly double.

Some actuaries refer to assets relative to payroll as the plan's "stability ratio." The higher this figure gets, the more variable plan costs become. For example, at a ratio of five, an actuarial loss of 10 percent in investment earnings represents a loss in assets of 50 percent of payroll. After smoothing, depending on amortization (let's assume a factor of 10), this results in a cost change of around 5 percent of pay. However, if the stability ratio reaches 10, a 10 percent loss becomes 100 percent of payroll, with a cost impact of 10 percent of pay. The sensitivity of the plan to market performance has doubled.

In the case of the plan in Graph 9, the board, staff, and stakeholders need to be aware of the increasing risk associated with their plan. The increasing cost volatility has implications for investment policy, funding methods, and employer budgeting, just to name a few areas.

Graph 9
Projection of assets as a percentage of payroll for a large county plan.
This projection assumes all actuarial assumptions are exactly met.
The horizontal axis is years. The vertical axis is plan assets as a percentage of pay.

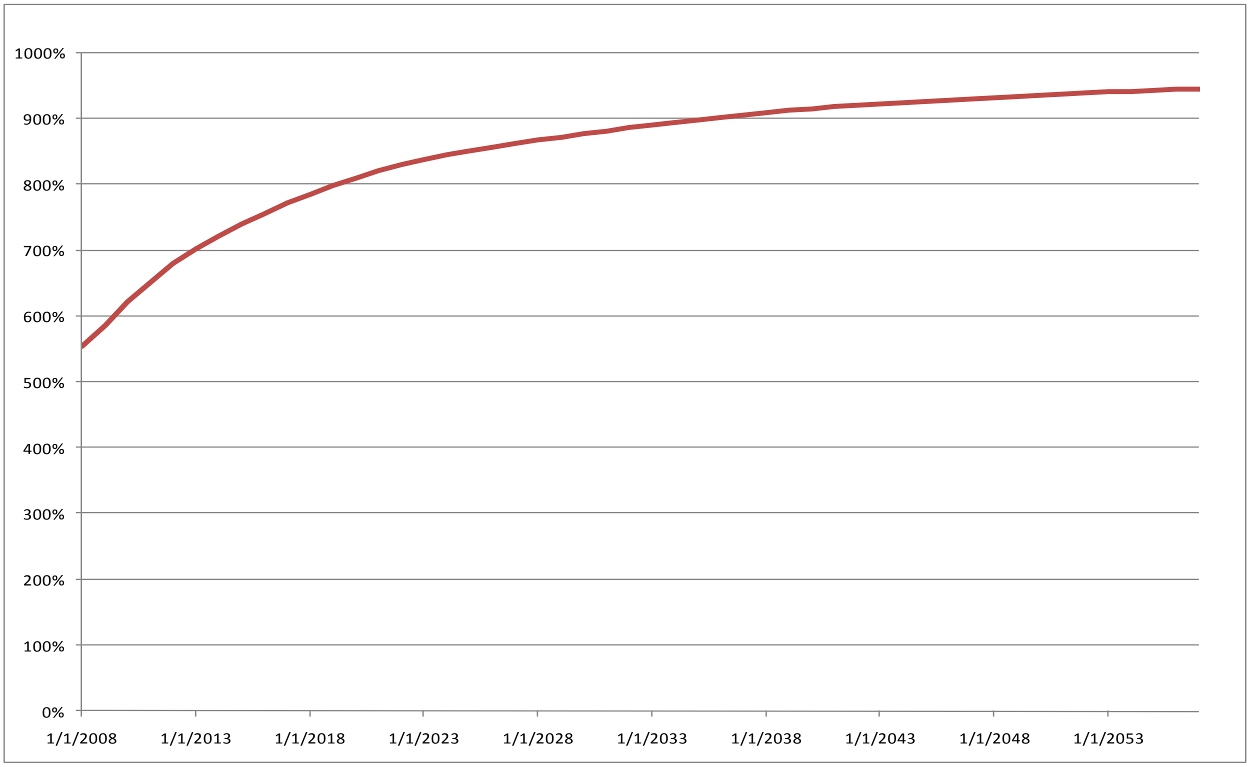

While simple projections of cost, funding, and assets are, well, simple, they can tell us a lot about the future of our plans. They should be a part of every major actuarial study.

In terms of our criteria for evaluation,

- Simplicity – Good
- Relevance – Fair

  While not a comprehensive as a simulation study, well-chosen projections can indicate plan dynamics that are noteworthy, important, and merit further study.

- Probative Value – Fair

  I find it odd that I rate a simple projection better in probative value than a range projection. However, I stand by my choice. It is the misleading nature of the range projection that I find of concern, the tendency for the eye to assume that all possibilities are between the extremes. As noted above, some very interesting observations can come from simple actuarial projections

- Consistency – Good

  Other than the starting point and the delaying effects of asset smoothing, there will generally be little change in projections from year to year.

- Cost – Good

  This is one of the most basic tools of every consulting actuary.

- Clarity – Good

  It is usually very clear what is being projected, although the full implications generally require some explanation and context.

Risk-Free Discount Rate

There has been a great deal of discussion of the desirability of disclosing the "market value" of pension liabilities. In this section we discuss this figure as a vehicle for communicating risk.

In our discussion here, we will define the market value of plan liabilities as the present value of the accrued benefit cash flows, with future payments discounted using a yield curve. Liabilities would be measured as an accumulated benefit obligation (ABO); present values could use a Treasury or corporate yield curve. We'll refer to this as the MVABO, for short.

To explore some of the ramifications of this measure, let's look at the example plan with the cash flows in Graph 10, below.

This plan has an actuarial accrued liability on a projected unit credit basis of $500 million, and it has about $400 million in assets. If we were to discount the ABO at a long-term government yield rate, the liability would be about $700 million. When the graph was prepared, a significant benefit improvement was being negotiated between the employer and its largest union.

Let's investigate how the market value of ABO measure behaves. To do this, I'm going to make a simplification: Rather than discounting the ABO cash flow with the entire yield curve, I'm going to use the long-term government bond yield. This is unlikely to make much difference in the results. When we compute the MVABO using the various long-term bond rates that have prevailed since 1926, we find values ranging from $275 million to $1.1 billion. However, more interesting than this wide range is the yearly *variation* in the measure.

In Graph 11, we look at the annual variation in the MVABO. For each year since 1926 we compute the MVABO for both the year in question and the following year; then we subtract the two. The same cash flow is used for each year. The distribution of the annual difference in the MVABO is shown in Graph 11.

Graph 10

Projection of accrued benefits in dollars for a transit district plan.
Benefits are shown both before (“old”) and after a plan improvement (“new”).
The horizontal axis is years. The vertical axis is plan benefits in dollars.

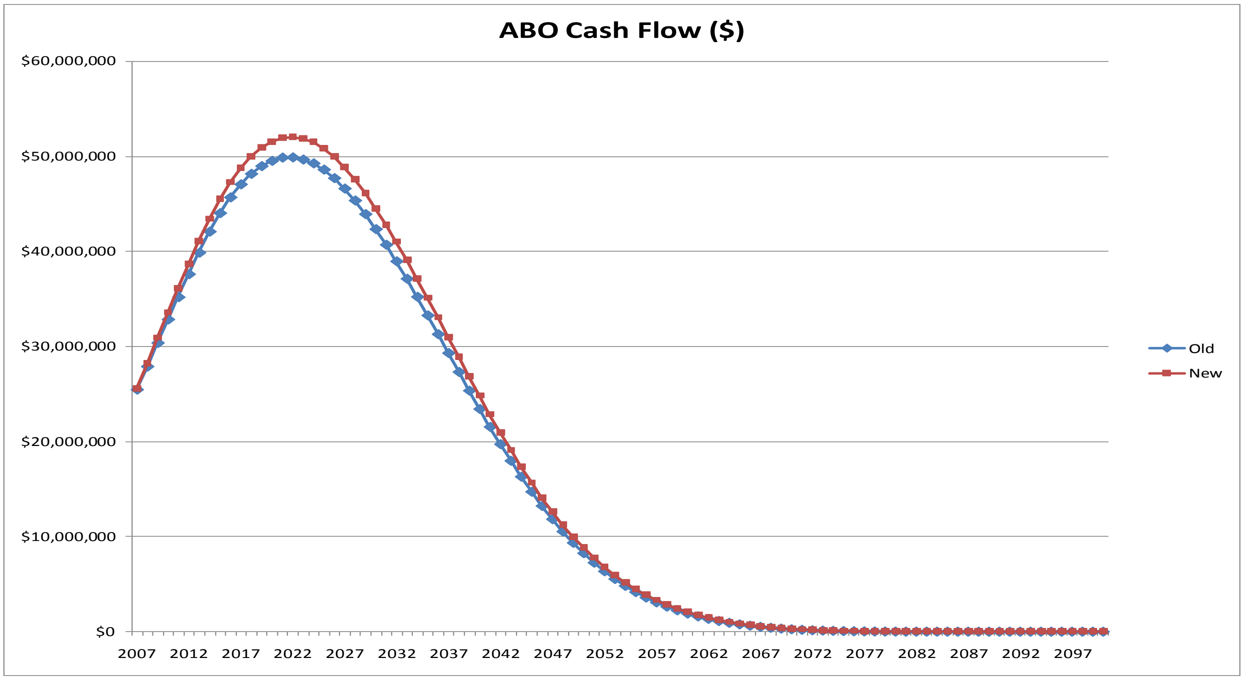

Graph 11

Distribution of the annual change in discounted ABO for a transit district plan.
Benefits are shown before plan improvement and are discounted at the long-term government bond yield rate. The horizontal axis is the change in discounted value.
The vertical axis is the number of years with that change.

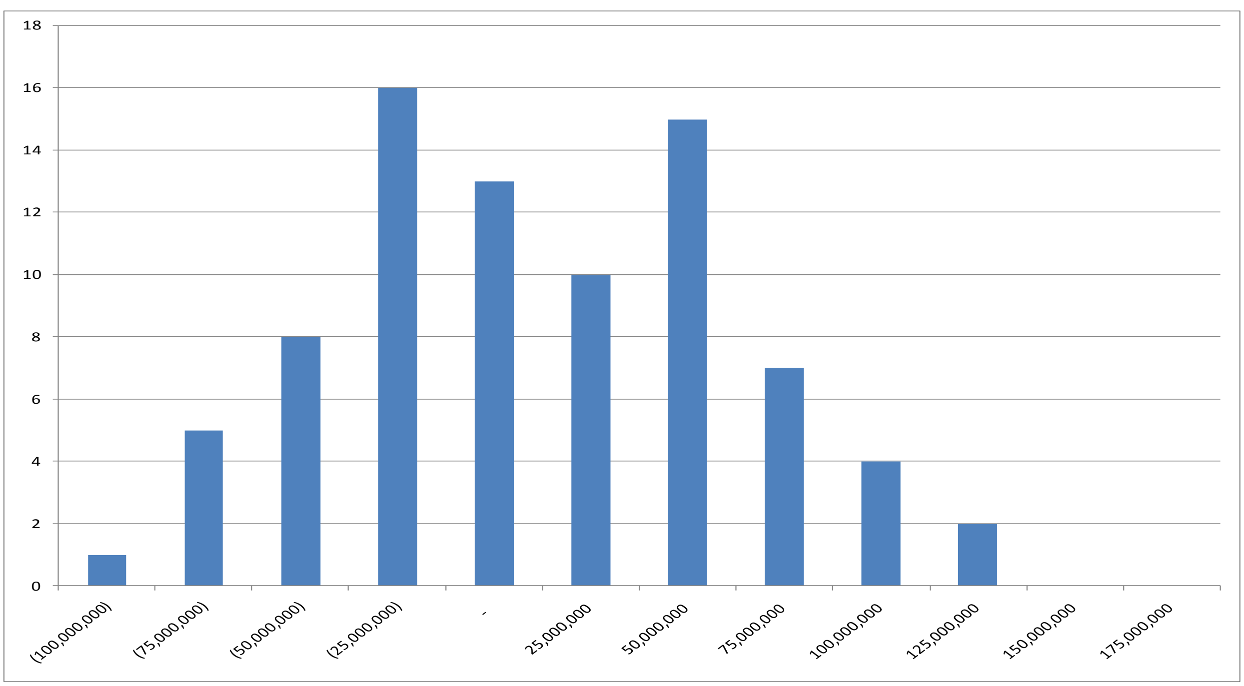

In the histogram in Graph 11 we see that the range of annual variation is very large: From -$100 million to $125 million. It's more or less centered at zero, which gives us some comfort. Is there a way to measure the yearly change relative to a standard? In fact, there is: We can compare the annual variation with the change in MVABO due to the plan amendment.

The average increase in the MVABO from the plan amendment across all years of Treasury yields was $24 million. This was actually a pretty significant benefit improvement; the modest increase of $24 million when compared with the $700 million average occurred because benefits for retired and inactive members were not changed.

If we compute the MVABO one year with the old benefits, and the next year with the new benefits, and then subtract the two, we get the distribution of differences shown in Graph 12.

Graph 12

Distribution of the annual change in discounted ABO for a transit district plan.
MVABO is shown before and after plan changes and is discounted at the long-term government bond rate.
The horizontal axis is the change in discounted value. The vertical axis is the number of years with that change.

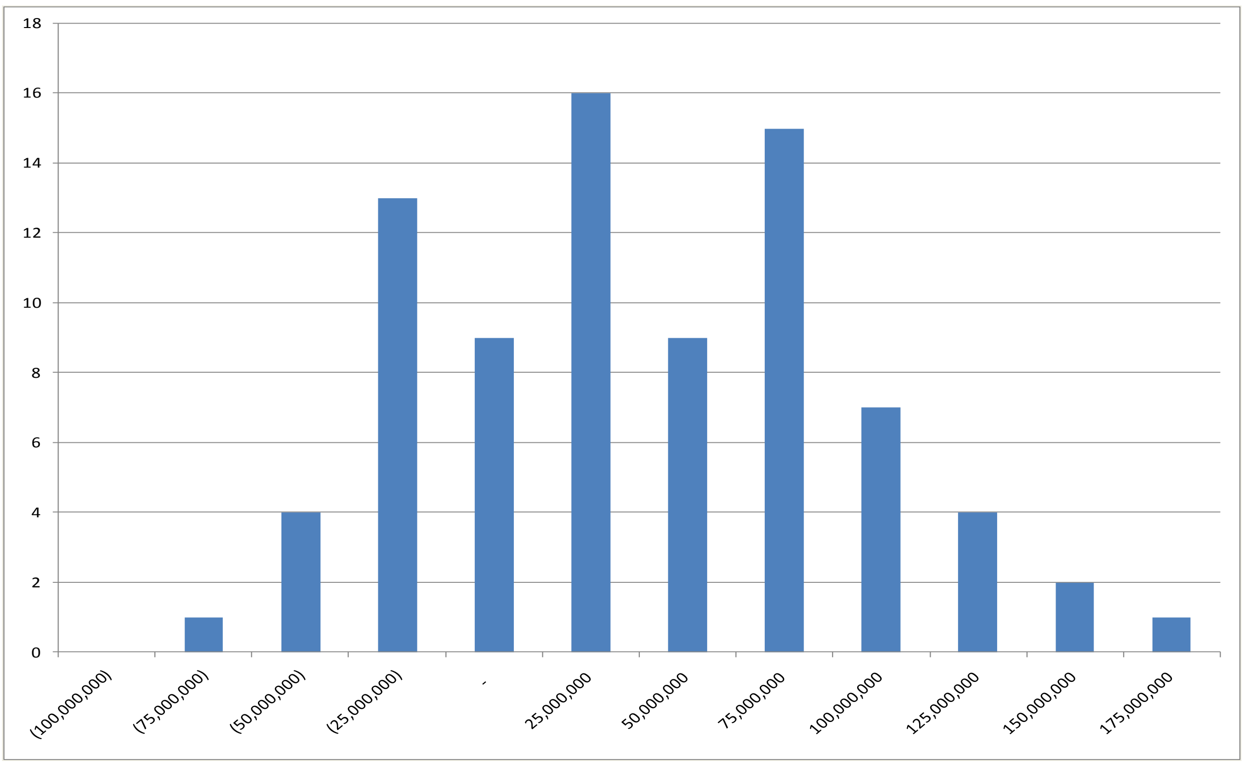

When we consider the benefit change, our histogram moves about $25 million to the right. The key factor to note is that fully one-third of the time, despite the improvement in plan benefits, the MVABO decreases from one year to the next. In short, changes in the discount rate frequently obscure even significant plan changes. It is the yield curve, rather than plan liabilities, that is being measured.

So how do we rank market value of liabilities on our rating system?

- Simplicity – Good

  Computation of cash flows and the related present value is trivially easy. Presentation of a single number could not be simpler.

- Relevance – Poor

  For an ongoing plan the MVABO is irrelevant. In most jurisdictions the liability being measured, the ABO, does not represent the legal obligation of the employer. Instead, the employer is legally responsible for the entire benefit formula for current active and inactive members, whether the benefits have accrued or not. Therefore, the MVABO does not measure the value of the promise.

  In addition, the MVABO does not measure the cost of plan benefits. Benefit cost, and taxes paid to provide them, will depend on the entire benefit formula and on the investment earnings from a diversified portfolio. The taxes levied will not depend on a notional portfolio of low risk bonds.

  MVABO does have relevance for terminating and frozen plans and for employers withdrawing from ongoing plans. In the latter case, if a participating employer wants to withdraw from a multiple-employer plan, with no possibility of contributing again, the MVABO should be considered as a possible termination liability.

- Probative Value – Poor

  It is impossible for a single number (however large) to even hint at some of the interesting and complicated probabilistic plan behavior we have discussed above.

- Consistency – Poor

  As noted above, the MVABO is strongly dependent on the level and shape of the yield curve, which changes daily. Accordingly, consistency is quite poor.

- Cost – Good

  This is another cheap spreadsheet exercise.

- Clarity – Poor

  The MVABO is actively misleading, particularly when combined with a comparison with plan assets.

  Actuarial funding accumulates assets using an asset target defined by the actuarial accrued liability, not based on the MVABO. Consequently, while assets can be meaningfully compared with the funding target, comparison with the (usually higher) MVABO presents an inaccurate picture: The funding method is simply not designed to accumulate assets equal to the MVABO.

Furthermore, the MVABO is often presented as the “real” or “true” liability of the plan, despite its manifest limitations. This frequently causes significant misinterpretation in the press, with a resulting loss in political credibility for the plan.

Conclusion

For those of you who haven't been keeping score at home, the following table summarizes our scoring of the risk communication techniques:

	Simplicity	Relevance	Probative Value	Consistency	Cost	Clarity
Full Simulation	Fair	Good	Good	Good	Poor	Good
Partial Simulation	Fair	Fair	Good	Good	Good	Good
Error Bars	Fair	Fair	Poor	Good	Good	Poor
Backtesting	Good	Fair	Poor	Poor	Good	Good
Range Projection	Good	Fair	Poor	Good	Good	Poor
Simple Projection	Good	Fair	Fair	Good	Good	Good
Risk-free Discount	Good	Poor	Poor	Poor	Good	Poor

There are a few additional conclusions that we can draw from our discussion:

1. Determinism is dead.

   Public pension plans are complex dynamic systems. They cannot be adequately studied with deterministic techniques. Our approach must change to be probabilistic, based on statistics and computer simulation.

2. Present values stink.

   Taking all or a portion of future cash flows, discounting them at one or more interest rates to a single point, and making decisions on that single number is insane. These systems are simply too complicated for that.

3. Liabilities are misleading.

   As noted above, liabilities are neither certain nor stable. They do not form a sufficient basis for determining or communicating plan risk. They are inadequate for informed decision making.

4. Stochastic techniques are fun.

   It is great fun and very rewarding to experiment with these plans using simulation models and then discuss the results with plan staff, boards and stakeholders. What is more important, those most responsible for determining the future of the plan are given the information they need for more informed decision-making.

It's been a pleasure writing this article, and I hope you enjoyed reading it. My advice to other actuaries would be to encourage you to develop a sense of playful exploration with regard to public plans. Build models, test, think, and test some more. Above all, don't rely on given wisdom to determine the best approach: Trust in yourselves.

Robert T. McCrory, FSA, is the executive vice president and director of the Northwest region of EFI Actuaries in Seattle, Wash. and can be reached at *BobMcCrory@qwest.net*.